

DescriptionField of the Invention

The present invention relates to a method for adjusting current clock counts for each computer site in a distributed computer system to effectively execute timestamped transactions.

Description of the Background

A distributed computer system comprises a number of computer sites, each of which has a processor and stores a database in a disk. The computer sites are interconnected by a network of communication lines. When a site receives a data processing request from a user, a processor in the site determines which database in the system includes the necessary data for processing by referring to a system catalog that stores information for data location. If the necessary data is stored in the database of the same computer site as the processor, the processor reads the necessary data from its database, executes data processing operations on the data, and writes the result into the database. A series of these processes is referred to as a transaction. If the necessary data is stored in a database at another computer site, the processor sends a request to read the necessary data from the database to the other computer site using the communication lines of the network. A processor of the other computer site reads the data from its database and sends the data to the computer site which sent the request using the communication lines. Then the processor of the computer site executes data processing operations on this data.

Figure 1 shows an example of two transactions which access the same data entry. It is assumed that the computer site corresponds to a branch of a bank, and the database stores data entries (records) representing the balance deposit for each of the bank's customers. When money is deposited or withdrawn from an account at any branch, a transaction must be performed by the computer system. During the transaction, the balance in the customer's account is read from a database, the money that is newly deposited or withdrawn by the customer or the bank is added to or subtracted from the balance for that customer, and the arithmetic result is written back into the data entry (record) in the database that corresponds to the account for that customer.

In this example, a data entry in the database at computer site A stores a customer's balance. It is assumed that the customer's beginning balance is \$50. At computer site A, an additional \$100 is deposited in the customer's account by a transaction that will be designated Ta. Immediately after, at another computer site B, \$30 is deposited in that same account by another transaction that will be designated Tb. Under these circumstances, a processor at computer site A will execute a READ operation for transaction Ta. The processor will add \$100 to

the beginning balance of \$50, and then writes \$150 into the database as the new balance in the account during a WRITE operation for transaction Ta. However, a processor at computer site B may send a request for a READ operation for transaction Tb to computer site A, and the processor at computer site A execute the READ operation before execution of the WRITE operation for transaction Ta. As a result, the processor at computer site B will read the customer's beginning balance of \$50. Then, the processor at computer site B adds \$30 to the beginning balance of \$50, and sends \$80 as a request for a WRITE operation to computer site A. The processor at computer site A writes \$80 into the database as the customer's new balance during a WRITE operation for transaction Tb. Therefore, the data entry in the database for this customer's account will indicate an ending balance of \$80. However, the ending balance in this customer's account is actually \$180. To prevent such errors, a concurrency control mechanism is used, i.e. a timestamp method. According to this method, each computer site has its own clock. When a request is made to initiate a transaction, the computer site assigns a timestamp to the new transaction. Preferably, the timestamp includes a transaction clock count that corresponds to a current clock count for the computer site at that time. For example, if a customer deposits \$100 at an automatic teller machine (ATM), a transaction request is generated by the computer site where the ATM is located in order to add \$100 to the customer's account. This computer site determines the computer site corresponding to the customer's account by examining a key recorded on the customer's cash card. The computer site also examines a current clock count for the site that is provided by its clock and clock register. The computer site will generate a transaction request, consisting of a operation (READ), a record key (100), and a table name (ACCOUNT), as shown in Figure 2.

Figures 3A and 3B show an example of two transactions that access the same data entry and timestamp table according to a timestamp method. It is assumed that computer site A and computer site B each have a clock and clock register that generate an accurate current clock count. When a transaction request is made by a customer or computer user at a branch corresponding to computer site A, computer site A assigns a transaction clock count and a computer site identifier to the transaction Ta. In this example, the current clock count is 10:00 and the identifier of the computer site is (A). As a result, computer site A assigns a timestamp 10:00 (A) to the transaction Ta, and will subsequently write this timestamp into the read timestamp field of the timestamp table corresponding to the customer's account when the transaction Ta reads data from the record. In this example, the write timestamp 0:00 is read out from timestamp table as shown in upper of Fig. 3B. The timestamp 10:00 (A) is larger than the write timestamp 0:00. Therefore, the transaction Ta reads data from the record and writes the timestamp 10:00 (A) in the read timestamp field to

the timestamp table as shown in middle of Fig. 3B.

Fig. 4 shows the data structure of a timestamp table and corresponding database (table account). The timestamp table is previously stored in memory different from database. The timestamp table consists of a record key, a read timestamp, and a write timestamp. The database consists of a key, a name, a data entry. In this example, the data entry is balance deposit because this database is table account. Each item of the timestamp table corresponds to item of the database according to the key. For example, Mike's balance deposit was read and written by a transaction generated by computer site A at 10:00. Immediately after computer site A requests a transaction, another transaction may be made at the branch corresponding to computer site B. Computer site B assigns the current timestamp 10:01 (B) to the new transaction Tb if the current clock count at computer site B is 10:01 when the transaction request is generated. When computer site B sends a request of a READ operation to computer site A and the computer site A reads the data entry by record key 100, the write timestamp 0:00 corresponding to the record key 100 is read out. The timestamp 10:01 (B) is larger than the write timestamp 0:00. Therefore, the transaction Tb reads data from the record and writes the timestamp 10:01 (B) in the read timestamp field of the timestamp table as shown in lower of Fig. 3B. When transaction Ta writes execution result to the record, the read timestamp 10:01 (B) is read out from the timestamp table. The timestamp 10:00 (A) of transaction Ta is smaller than the read timestamp 10:01 (B). Therefore, write processing of transaction Ta is rejected. Then, transaction Ta is aborted in computer site A. When transaction Tb writes execution result to the record, the read timestamp 10:01 (B) is read out from the timestamp table. The timestamp 10:01 (B) of transaction Tb is not smaller than the read timestamp 10:01 (B). Therefore, write processing of transaction Tb is accepted. Then, transaction Tb is committed (succeeded) in computer site B. (Transaction Ta is assigned larger timestamp and executed again.)

Unfortunately, the current clock counts in different computer sites are not always accurate and are not always the same. Even if the current clock counts in different computer sites are identical same during an initialization mode, as time progresses, the difference in the current clock counts of two computer sites will become larger and larger.

Figures 5A and 5B show an example of the attempted execution of two transactions by two computer sites that have different current clock counts and timestamp table. In this example, the current clock count of computer site A is early when compared to the current clock count of computer site B, and the current clock count of computer site B is late or delayed when compared to the current clock count of computer site A.

In this example, when a transaction request Ta is made at the branch corresponding to computer site A, computer site A assigns a transaction clock count 10:00

and an identifier (A) to the transaction request. Preferably, transaction clock count 10:00 corresponds to the current clock count at computer site A. Immediately thereafter, a transaction request Tb is made at the branch corresponding to computer site B, and computer site B assigns a transaction clock count 9:55 and an identifier (B) to the transaction request. The transaction clock count 9:55 corresponds to the current clock count at computer site B, which is late when compared with the current clock count at computer site A.

When computer site A executed a READ operation according to the transaction Ta, timestamp 10:00 (A) is written in read timestamp field of timestamp table as shown in upper of Fig. 5B. This processing is the same as that shown in Fig. 3A. When computer site A executes a READ operation according to the transaction Tb, write timestamp 0:00 is read out from timestamp table. The timestamp 9:55 (B) is larger than the write timestamp 0:00. Therefore, computer site A read the data entry by record key 100 and write the timestamp 9:55 (B) in read timestamp field as shown in the middle of Fig. 5B. When computer site A writes execution result of transaction Ta to the record, the read timestamp 9:55 (B) is read out from timestamp table. The timestamp 10:00 (A) is larger than the read timestamp 9:55 (B). Therefore, computer site A executes a WRITE operation and writes the timestamp 10:00 (A) in write timestamp field of timestamp table as shown in the lower of Fig. 5B. Then, transaction Ta is committed in computer site A as shown in Fig. 5A. When computer site A writes execution result of transaction Tb to the record, the read timestamp 9:55 (B) and write timestamp 10:00 (A) is read out from timestamp table. The timestamp 9:55 (B) is not smaller than the read timestamp 9:55 (B), but it is smaller than the write timestamp 10:00 (A). Therefore, write processing of transaction Tb is rejected. After transaction Tb is aborted, the current clock count for computer site B may be 10:01. In this situation, it is preferable for computer site B to write a new transaction clock count of 10:01 into the data structure corresponding to transaction request Tb. As a result, a new timestamp 10:01 (B) is assigned to transaction Tb.

Now, if computer site A again attempts to execute a READ operation according to transaction Tb, transaction Tb will not be aborted because the write timestamp 10:00 (A) is smaller timestamp 10:01 (B) assigned to transaction Tb. See Figure 3.

Thus, if one computer site executing a new transaction attempts to execute a READ operation on the database at a particular record key, but the read/write timestamp corresponding to the data entry is later than the timestamp assigned to the new transaction, then the computer site must abort the new transaction. As the difference in current clock counts between different computer sites becomes larger, the number of transactions that must be aborted by the computer site with the late current clock count increases.

To address this problem, a method for adjusting cur-

rent clock counts is disclosed in Japanese Patent Disclosure (Kokai) P59-55553. Figure 6 shows a distributed system computer taught by this reference. In Figure 6, three computer sites 60, 62, and 64 already are interconnected by a network, and a new computer site 66 is being connected in the network.

As disclosed in Japanese Patent Disclosure (Kokai) P59-55553, new computer site 66 inquires about the current clock counts provided by the clocks and clock registers in existing computer sites 60, 62, and 64. When new computer site 66 receives information indicating the current clock counts of each of the existing computer sites, a clock offset OFS is calculated. The clock offset OFS is the amount of time that the current clock count corresponding to new computer site 66 will be advanced (i.e., the clock is made early rather than late).

In a distributed computer system of the type disclosed in Japanese Patent Disclosure (Kokai) P59-55553, the clock offset OFS is calculated by new computer site 66 by first selecting the current clock count from all of the computer sites that corresponds to the clock that is the fastest (e.g., 10:05 is selected instead of 10:00). Then, a clock offset OFS is added to the current clock count of new computer site 66 that will make its current clock count identical to the current clock count of the computer site with the fastest (i.e., earliest) clock. This process of calculating a clock offset OFS is repeated whenever a new computer site is connected to the network.

However, when using the method disclosed in Japanese Patent Disclosure (Kokai) P59-55553, it is difficult to accurately adjust the current clock count of the new computer site. Information indicating the current clock counts of each of the existing computer sites is received by new computer site 66 only after a communication delay, and this communication delay must be estimated correctly to determine the amount of time that the current clock count corresponding to new computer site 66 should be advanced.

Moreover, the method disclosed in Japanese Patent Disclosure (Kokai) P59-55553 does not provide any method for adjusting the current clock counts of existing computer sites. As time passes, the differences between the current clock counts of these computer sites will increase.

SUMMARY OF THE INVENTION

It is an object of the invention to overcome at least some of the problems associated with the prior art.

It is therefore desirable to provide an improved method for adjusting current clock counts in a distributed computer system. It would be beneficial to provide a process in which the current clock count corresponding to one computer site is changed to more closely match the current clock count corresponding to another computer site.

Additional objects and advantages of the invention

will be set forth in part in the description which follows, and in part will be obvious from the description, or may be learned by practice of the invention. The objects and advantages of the invention may be realized and obtained by means of the instrumentalities and combinations particularly pointed out in claims 1 and 10 and in the dependent claims.

To achieve the objects and in accordance with the purposes of the invention, as embodied and broadly described herein, a method is provided for adjusting current clock counts in a distributed computer system. The distributed computer system includes a number of computer sites, each computer site including a processor, a memory having a number of record keys for storing timestamp table, a database, an abort register for storing an abort count corresponding to the computer site, a clock for generating a clock signal indicating the passage of time, and a clock register responsive to the clock signal for storing a current clock count corresponding to the computer site. The database for each computer site includes a number of data entries (records) that correspond to the record key and that can be accessed by a processor. Each record key corresponds to timestamp which includes at least clock count that indicates when the corresponding data entry in the corresponding record key was most recently accessed. The distributed computer system also includes a network for interconnecting the computer sites to enable a processor in one computer site to send the data entries and request for accessing in another computer site.

The method includes the following steps, which are executed by a processor in at least a first one of the computer sites: assigning a timestamp to a new transaction that is related to the current clock count for the first computer site, at times when the processor in the first computer site is initiating the new transaction; reading the timestamp corresponding to the record key from the memory, at times when the processor in the first computer site is initiating the new transaction; comparing an existing timestamp corresponding to the record key with the timestamp assigned to the new transaction; aborting the new transaction, at times when the timestamp assigned to the new transaction is smaller than the existing timestamp corresponding to the record key; incrementing the abort count corresponding to the first computer site, at times when the new transaction is aborted, to indicate the number of times that new transactions initiated by the first computer site have been aborted; accessing an abort count corresponding to a second computer site; comparing the abort count corresponding to the first computer site with the abort count corresponding to the second computer site; and changing the current clock count corresponding to the first computer site to more closely match the current clock count corresponding to the second computer site, at times when the comparison between the abort counts in the first and second computer sites indicates that new transactions have been aborted more frequently in one of the two computer sites.

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate preferred embodiments of the invention and, together with the description of the invention, explain the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows a prior art example of the execution of two transactions involving the same database.

Figure 2 shows the data structure of a transaction request.

Figures 3A and 3B show a prior art example of the execution of two transactions involving the same database and using timestamps.

Figure 4 shows the data structure of a timestamp table and a database.

Figures 5A and 5B show a prior art example of the execution of two transactions involving the same database and using timestamps.

Figure 6 shows a prior art example of a distributed computer system.

Figure 7 shows a block diagram of computer sites in a distributed computer system.

Figure 8 shows the data structure of a timestamp.

Figure 9A and 9B show a flow chart of a method for adjusting current clock counts in a distributed computer system.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 7 shows a distributed computer system. The distributed computer system includes three computer sites 1, 2, and 3. The computer sites are interconnected by a network 4.

Computer sites 1, 2, and 3 respectively include input devices 11, 21, and 31, which can be operated by a computer site user or a customer to request that a new transaction be initiated, memories 12, 22, and 32 for storing timestamp table and system catalog, processors 13, 23, and 33 for executing transactions, abort registers 14, 24, and 34 for storing an abort count corresponding to the computer site, clocks 15, 25, and 35 for generating a clock signal indicating the passage of time, clock registers 16, 26, and 36 that are responsive to the clock signal for storing a current clock count corresponding to the computer site, and databases 17, 27, and 37 for storing a plurality of data entry.

When the distributed computer system is in opera-

tion, a customer or computer user located at computer site 1 will request that a new transaction be initiated using input device 11. Processor 13 in computer site 1 is coupled to input device 11, and responds to the request by generating a transaction request that will be designated T1 to distinguish it from transaction requests generated by processors 23 and 33 in computer sites 2 and 3. The transaction request includes an operation (e.g., add, subtract), a record key, and a data (e.g., an amount to be deposited or withdrawn from a checking account). This information contained in the transaction request is typically supplied to processor 13 by the customer or computer user located at computer site 1 that requested initiation of new transaction T1.

In accordance with the invention, processor 13 assigns a timestamp to the new transaction T1 that is related to the current clock count for computer site 1, at times when processor 13 in computer site 1 is initiating the new transaction. Preferably, processor 13 assigns a timestamp $t_s(T1)$ to transaction T1 that has a transaction clock count that matches the current clock count stored in clock register 16 of computer site 1. As embodied herein, the timestamp also includes a site identifier. As a result, even if processors at two different computer sites generate separate transaction requests at the same time, the timestamps assigned to any two transactions will be different.

Figure 8 shows the data structure of a timestamp. The timestamp consists of " α " most significant bits for storing time when transaction is generated, and " β " least significant bits for storing an identifier that indicates in which site the transaction is generated. " β " must be large enough so that 2^β is not less than the total number of computer sites.

In accordance with the invention, a processor in a computer site will read the timestamp from the memory according to the record key, at times when the processor is initiating a new transaction. Next, the processor will compare the timestamp corresponding to the record key with the timestamp assigned to the new transaction. The processor will access the data entry (record) from the database according to the record key, at times when the timestamp assigned to the new transaction is larger than the timestamp corresponding to the record key. This processing is already explained in the Description of the Background.

As an example of the operation of a preferred embodiment of the invention, it is assumed that input device 21 requests that a new transaction T2 be initiated by processor 23 in computer site 2 almost the same time that transaction T1 is being initiated by processor 13 in computer site 1. Processor 23 generates a transaction request T2 and assigns a timestamp $t_s(T2)$ to transaction T2 in accordance with the current clock count stored in clock register 26 and the identifier of computer site 2.

In this example, it is assumed that transaction T1 and transaction T2 access the same data entry in the same database, i.e., the database 37 in computer site 3,

and that clock 15 and clock register 16 are late when compared with clock 25 and clock register 26. As a result, the current clock count assigned to transaction requests generated by processor 13 in computer site 1 will be late (e.g., 9:55 instead of 10:00) in comparison with the current clock count assigned to transaction requests generated by processor 23 in computer site 2.

Under these circumstances, there are four possible orders in which transactions T1 and T2 can be generated, and in which a data entry in the database of computer site 3 can be accessed while performing transactions T1 and T2:

(1) The actual time at which transaction T1 is generated by processor 13 in computer site 1 is before the actual time at which transaction T2 is generated by processor 23 in computer site 2; and

the data entry in the database of computer site 3 is accessed during transaction T1 before the data entry is accessed during transaction T2.

(2) The actual time at which transaction T1 is generated by processor 13 in computer site 1 is before the actual time at which transaction T2 is generated by processor 23 in computer site 2; and

the data entry in the database of computer site 3 is accessed during transaction T2 before the data entry is accessed during transaction T1.

(3) The actual time at which transaction T2 is generated by processor 23 in computer site 2 is before the actual time at which transaction T1 is generated by processor 13 in computer site 1; and

the data entry in the database of computer site 3 is accessed during transaction T1 before the data entry is accessed during transaction T2.

(4) The actual time at which transaction T2 is generated by processor 23 in computer site 2 is before the actual time at which transaction T1 is generated by processor 13 in computer site 1; and

the data entry in the database of computer site 3 is accessed during transaction T2 before the data entry is accessed during transaction T1.

In the first two examples above, because the current clock count in computer site 1 is later than the current clock count in computer site 2, and because the actual time at which transaction T1 is generated by processor 13 is before the actual time at which transaction T2 is generated by processor 23, the timestamp $t_s(T1)$ assigned to transaction T1 will be earlier than the timestamp $t_s(T2)$ assigned to transaction T2.

In a preferred embodiment of the invention, a processor in a first computer site will execute a new transaction, at times when the timestamp assigned to the new transaction is larger than the timestamp corresponding to the record key. Therefore, in the first example, proc-

essor 23 in computer site 2 executes a READ operation during transaction T2.

In accordance with the invention, a processor in a first computer site will abort a new transaction, at times when the timestamp assigned to the new transaction is smaller than the timestamp corresponding to the data. Therefore, in the second example, processor 13 in computer site 1 aborts transaction T1.

In the third and fourth examples above, even though the timestamp assigned to transaction requests generated by processor 13 is smaller than the timestamp assigned to transaction requests generated by processor 23, if the difference in the actual times at which the transactions are generated is greater than the difference in timestamps, then the timestamp $t_s(T2)$ assigned to transaction T2 will be smaller than the timestamp $t_s(T1)$ assigned to transaction T1.

Therefore, in the third example, processor 23 in computer site 2 aborts transaction T2. In the fourth example, processor 13 in computer site 1 executes a READ operation during transaction T1.

On the other hand, the timestamp $t_s(T1)$ assigned to transaction T1 will be smaller than the timestamp $t_s(T2)$ assigned to transaction T2 if the difference in the actual times at which the transactions are generated is smaller than the disparity in timestamps.

Under these circumstances, in the third example, processor 23 in computer site 2 executes a READ operation during transaction T2. In the fourth example, processor 13 in computer site 1 aborts transaction T1.

In accordance with the invention, a processor in a first computer site will change the abort count corresponding to the first computer site, at times when a new transaction is aborted, to indicate the number of times that new transactions initiated by the first computer site have been aborted. Preferably, the abort count for computer site 1 is stored in abort register 14. The abort counts corresponding to computer sites 1, 2, and 3 are continually updated by processors 13, 23, and 33 by the incrementing of the abort counts stored in abort registers 14, 24, and 34, respectively, whenever processors 13, 23, and 33 abort a transaction.

The probability that transaction T1 will be aborted is larger than the probability that transaction T2 will be aborted in the examples above because the current clock count of the timestamp assigned to transaction requests generated by processor 13 is late or delayed (e.g., 9:55 instead of 10:00) in comparison with the current clock count of the timestamp assigned to transaction requests generated by processor 23. As a result, a larger number of transactions will be aborted by computer sites that have a clock and/or clock register that are late when compared with the clocks and clock registers for the other computer sites in the system.

In the present invention, the current clock counts in a distributed computer system are adjusted. In accordance with the invention, a processor in a first computer site accesses an abort count corresponding to a second

computer site. The processor then compares the abort count corresponding to the first computer site with the abort count corresponding to the second computer site. At times when the comparison between the abort counts in the first and second computer sites indicates that new transactions have been aborted more frequently in one of the two computer sites, the processor changes the current clock count corresponding to the first computer site to more closely match the current clock count corresponding to the second computer site.

Preferably, the computer sites recurrently executes the steps required to adjust their current clock counts. For example, a processor may perform these steps after it has not executed any transactions for a predetermined time. Alternatively, a user clock matching request can be entered at a computer site using the corresponding input device, and the computer site then executes the steps required to adjust current clock counts in response to entry of the user clock matching request.

Figure 9A and 9B are a flow chart showing a preferred embodiment of a method for adjusting current clock counts in a distributed computer system.

First, an interruption is supplied to processor 13 in computer site 1 in step 91. In step 93, processor 13 determines the number of times it has aborted a new transaction by reading abort register 14 in computer site 1. Then, in step 93, processor 13 temporarily defines the parameter MIN as equal to n_1 , the number of transactions aborted by processor 13. The parameter MIN is then sent to the next site, i.e., computer site 2. Generally speaking, the parameter MIN will represent the minimum number of transactions that have been aborted by a single one of the computer sites in the distributed computer system.

At step 95, processor 23 determines the number of times (n_2) it has aborted a new transaction by reading abort register 24 in computer site 2. Processor 23 then compares n_2 with MIN. If the number n_2 of transactions aborted by processor 23 is smaller than the value of MIN received by processor 23 from computer site 1, then processor 23 in step 97 changes the value of MIN to equal n_2 , the number of transactions aborted by processor 23.

During step 99, the identifier of the computer site that most recently received and processed the parameter MIN is compared with total number of computer sites (m) to determine if all computer sites in the computer system have had an opportunity to update the value of MIN. If some computer sites still have not received the parameter MIN for processing, the current processor (e.g., processor 23) sends the parameter MIN to the next site, i.e., computer site 3, in step 101.

After each computer site has received and processed the parameter MIN, the current processor (e.g., processor 33) sends the parameter MIN back to the original computer site, e.g., computer site i, in step 103. At step 105, processor 13 in computer site i compares the number n_i of transactions aborted by it with the parameter

MIN, which now corresponds to the minimum number of transactions that have been aborted by any single the computer site in the distributed computer system.

5 If the difference between the number n_i of transactions aborted by processor i3 is larger than the value of MIN by at least a threshold value δ , a parameter g is calculated in step 107. The value of g is a function of the size of the difference between n_i and MIN. In step 109, 10 processor i3 in computer site i moves its current clock count forward by an amount equal to the value of g.

At step 111, processor i3 in computer site i determines whether all of the computer sites in the system have had an opportunity to adjust their current clock counts. If not all of the computer sites have executed step 105, the parameter MIN is then sent to the processor in one of these other computer sites in step 113. After all computer sites have executed step 105, the process for adjusting current clock counts has been completed.

20 In the preferred embodiment of the invention, abort counts are stored in abort registers 14, 24, and 34. In another embodiment, the abort registers can be replaced by succeed registers that store succeed counts.

In this second embodiment, a processor in a first computer site will change the succeed count corresponding to the first computer site, at times when a new transaction is not aborted, to indicate the number of times that new transactions initiated by the first computer site have been successfully completed. The succeed counts corresponding to computer sites 1, 2, and 3 would be continually updated by processors 13, 23, and 33 by the incrementing of the succeed counts stored in the succeed registers whenever processors 13, 23, and 33 successfully complete a transaction.

35 In the second embodiment, a processor in a first computer site accesses an succeed count corresponding to a second computer site. The processor then compares the succeed count corresponding to the first computer site with the succeed count corresponding to the second computer site. At times when the comparison between the succeed counts in the first and second computer sites indicates that new transactions have been successfully completed more frequently in one of the two computer sites, the processor changes the current clock count corresponding to the first computer site to more closely match the current clock count corresponding to the second computer site.

40 In the second embodiment, a process similar to that shown in Figures 9A and 9B is carried out, but instead of the parameter MIN, a parameter MAX is defined that will represent the maximum number of transactions that have been successfully completed by a single one of the computer sites in the distributed computer system. The value of g for this embodiment is a function of the size of the difference between n_i , the number of transactions succeeded by processor i3, and MAX. Processor i3 in computer site i will move its current clock count forward by an amount equal to the value of g at times when the

difference between n_i and MAX is larger than a threshold value of δ .

Preferably, processor I3 in computer site i will clear its abort count or succeed count whenever processor I3 has just adjusted its current clock count. In this case, the difference in the value of the abort or succeed counts between various computer sites will more directly reflect the amount of adjustment that must be made in current clock counts.

The function used to calculate g is represented as a (constant) x (average generation interval for transaction requests) \times $(n_i - MIN)/MIN$. The term $(n_i - MIN)/MIN$ serves to dynamically change the current clock count in a computer site which has a large number of abort of transaction. The use of the average generation intervals for transaction requests serves to change current clock counts in proportion to the average generation interval, i.e., if this average is short, it indicates that many transactions are being executed at almost the same time. Therefore, it is necessary for current clock counts to change minutely. The constant used to calculate g should be less than one, in order to gradually improve any unbalance in transaction processing of each computer site.

Claims

1. A method for adjusting current clock counts in a distributed computer system, wherein the distributed computer system includes a plurality of computer sites (1, 2, 3), each computer site including a processor (13), a memory (12) having a plurality of record keys for storing a timestamp table, a database (17), abort register means (14) for storing an abort count corresponding to the computer site, clock means (15) for generating a clock signal indicating the passage of time, and clock register (16) means responsive to the clock signal for storing a current clock count corresponding to the computer site, wherein the database for each computer site includes a plurality of data entries that correspond to the plurality of record keys and that can be accessed by a processor, wherein each record key corresponds to timestamp which includes at least clock count that indicates when the corresponding data entry in the corresponding record key was most recently accessed, wherein the distributed computer system includes network means (4) for interconnecting the plurality of computer sites to enable a processor in one computer site to send the data entries and request for accessing to another computer site, and wherein the following steps are executed by a processor in at least a first one of the computer sites:

assigning a timestamp to a new transaction that is related to the current clock count for the first computer site, at times when the processor in

5

10

15

20

25

30

35

40

45

50

55

55

the first computer site is initiating the new transaction;

reading the timestamp corresponding to the record key from the memory, at times when the processor in the first computer site is initiating the new transaction;

comparing an existing timestamp corresponding to the record key with the timestamp assigned to the new transaction;

aborting the new transaction, at times when the timestamp assigned to the new transaction is smaller than the existing timestamp corresponding to the record key, characterised by incrementing the abort count corresponding to the first computer site, at times when the new transaction is aborted, to indicate the number of times that new transactions initiated by the first computer site have been aborted;

accessing an abort count corresponding to a second computer site;

comparing the abort count corresponding to the first computer site with the abort count corresponding to the second computer site; and

changing the current clock count corresponding to the first computer site to more closely match the current clock count corresponding to the second computer site, at times when the comparison between the abort counts in the first and second computer sites indicates that new transactions have been aborted more frequently in one of the two computer sites.

2. A method in accordance with claim 1, in which:

the step of changing the current clock count comprises changing the current clock count by an amount that is proportional to the difference between the number of times that new transactions have been aborted in the two computer sites.

3. A method in accordance with claim 1, and further comprising:

changing the abort count corresponding to the first computer site to more closely match the abort count corresponding to the second computer site, at times when the current clock count corresponding to the first computer site is changed to more closely match the current clock count corresponding to the second computer site.

4. A method in accordance with claim 1, in which:

the step of changing the current clock count comprises changing the current clock count only if the number of new transactions aborted

in the first computer site exceeds the number of new transactions aborted in the second computer site.

5. A method in accordance with claim 1, in which: 5
the step of changing the current clock count comprises changing the current clock count only if the number of new transactions aborted in the first computer site exceeds the number of new transactions aborted in the second computer site by a threshold amount. 10

6. A method in accordance with claim 1, and further comprising: 15
entering a new read clock count into the timestamp corresponding to the record key to indicate when the corresponding data entry was accessed by the processor, wherein the new read clock count is related to the current clock count corresponding to the first computer site. 20

7. A method in accordance with claim 6, wherein each timestamp includes an identifier corresponding to each read clock count that indicates which processor accessed the corresponding data entry in the corresponding record key, and further comprising: 25
entering an identifier into the timestamp corresponding to the record key to indicate that the corresponding data entry was accessed by the processor in the first computer site, at times when a new read clock count is entered into the timestamp. 30
35

8. A method in accordance with claim 1, in which the first computer site recurrently executes the steps of accessing the abort count for the second computer site, comparing the abort count for the first computer site with the abort count for the second computer site, and changing the current clock count for the first computer site to more closely match the current clock count for the second computer site. 40
45

9. A method in accordance with claim 1, in which the first computer site includes input means (11) for entering a user clock matching request, and in which the first computer site executes the steps of accessing the abort count for the second computer site, comparing the abort count for the first computer site with the abort count for the second computer site, and changing the current clock count for the first computer site to more closely match the current clock count for the second computer site, in response to entry of a user clock matching request. 50
55

10. A method for adjusting current clock counts in a distributed computer system, wherein the distributed computer system includes a plurality of computer sites (1, 2, 3), each computer site including a processor (13), a memory (12) having a plurality of record keys for storing a timestamp table, a database (17), succeed register means for storing an succeed count corresponding to the computer site, clock means (15) for generating clock signal indicating the passage of time, and clock register means (16) responsive to the clock signal for storing a current clock count corresponding to the computer site, wherein the database for each computer site includes a plurality of data entries that correspond to the plurality of record keys and that can be accessed by a processor, wherein each record key corresponds to timestamp which includes at least clock count that indicates when the corresponding data entry in the corresponding record key was most recently accessed, wherein the distributed computer system includes network means (14) for interconnecting the plurality of computer sites to enable a processor in one computer site to send the data entries and request for accessing to another computer site, and wherein the following steps are executed by a processor in at least a first one of the computer sites: 60
65
70
75
80
85
90
95
100
105
110
115
120
125
130
135
140
145
150
155
160
165
170
175
180
185
190
195
200
205
210
215
220
225
230
235
240
245
250
255
260
265
270
275
280
285
290
295
300
305
310
315
320
325
330
335
340
345
350
355
360
365
370
375
380
385
390
395
400
405
410
415
420
425
430
435
440
445
450
455
460
465
470
475
480
485
490
495
500
505
510
515
520
525
530
535
540
545
550
555
560
565
570
575
580
585
590
595
600
605
610
615
620
625
630
635
640
645
650
655
660
665
670
675
680
685
690
695
700
705
710
715
720
725
730
735
740
745
750
755
760
765
770
775
780
785
790
795
800
805
810
815
820
825
830
835
840
845
850
855
860
865
870
875
880

parison between the succeed counts in the first and second computer sites indicates that new transactions have been successfully completed more frequently in one of the two computer sites. 5

11. A method in accordance with claim 10, in which: the step of changing the current clock count comprises changing the current clock count by an amount that is proportional to the difference between the number of times that new transactions have been successfully completed in the two computer sites. 10

12. A method in accordance with claim 10, and further comprising: changing the succeed count corresponding to the first computer site to more closely match the succeed count corresponding to the second computer site, at times when the current clock count corresponding to the first computer site is changed to more closely match the current clock count corresponding to the second computer site. 15

13. A method in accordance with claim 10, in which: the step of changing the current clock count comprises changing the current clock count only if the number of new transactions successfully completed in the first computer site exceeds the number of new transactions successfully completed in the second computer site. 20

14. A method in accordance with claim 10, in which: the step of changing the current clock count comprises changing the current clock count only if the number of new transactions successfully completed in the first computer site exceeds the number of new transactions successfully completed in the second computer site by a threshold amount. 25

15. A method in accordance with claim 10, and further comprising: entering a new read clock count into the timestamp corresponding to the record key to indicate when the corresponding data entry was accessed by the processor, wherein the new read clock count is related to the current clock count corresponding to the first computer site. 30

16. A method in accordance with claim 10, wherein each timestamp includes an identifier corresponding to each read clock count that indicates which processor accessed the corresponding data entry in the corresponding record key, and further comprising: entering an identifier into the timestamp corresponding to the record key to indicate that the corresponding data entry was accessed by the processor in the first computer site, at times when a new read clock count is entered into the timestamp. 35

17. A method in accordance with claim 10, in which the time computer site recurrently executes the steps of accessing the succeed count for the second computer site, comparing the succeed count for the first computer site with the succeed count for the second computer site, and changing the current clock count for the first computer site to more closely match the current clock count for the second computer site. 40

18. A method in accordance with claim 10, in which the first computer site includes input means (11) for entering a user clock matching request, and in which the first computer site executes the steps of accessing the succeed count for the second computer site, comparing the succeed count for the first computer site with the succeed count for the second computer site, and changing the current clock count for the first computer site to more closely match the current clock count for the second computer site, in response to entry of a user clock matching request. 45

Patentansprüche

1. Verfahren, um die gegenwärtigen Zeittakte in einem verteilten Rechnersystem einzustellen, wobei das verteilte Rechnersystem eine Mehrzahl von Rechnerplätzen (1, 2, 3) und jeder Rechnerplatz einen Prozessor (13), einen Speicher (12) mit einer Mehrzahl von Aufzeichnungsschlüsseln zur Speicherung einer Zeitmarkentabelle, eine Datenbank (17), Abbruchregistermittel (14) zur Speicherung eines Abbruchzählers für den Rechnerplatz, Taktmittel (15) zur Erzeugung eines Taktsignals, das den Zeitablauf anzeigt, und ein Taktregister (16) aufweist, das auf das Taktsignal zur Speicherung eines gegenwärtigen Zeittakts für den Rechnerplatz reagiert, wobei die Datenbank jedes Rechnerplatzes eine Mehrzahl von Dateneinträgen aufweist, die der Mehrzahl von Aufzeichnungsschlüsseln entsprechen und auf die mittels eines Prozessors zugegriffen werden kann, wobei jeder Aufzeichnungsschlüssel einer Zeitmarke entspricht, die mindestens einen Zeittakt beinhaltet, der anzeigt, wann auf den entsprechenden Dateneintrag unter dem entsprechenden Aufzeichnungsschlüssel zuletzt zugegriffen

wurde, wobei das verteilte Rechnersystem Netzwerkmittel (4) zur Verbindung der Mehrzahl von Rechenplätzen aufweist, um es einem Prozessor in einem Rechnerplatz zu ermöglichen, Dateneinträge zu senden und den Zugriff auf andere Rechnerplätze anzufordern, und wobei die folgenden Schritte durch einen Prozessor mindestens in dem ersten der Rechnerplätze ausgeführt werden:

10 Zuteilung einer auf den gegenwärtigen Zeittakt des ersten Rechnerplatzes bezogenen Zeitmarke an eine neue Transaktion dann, wenn der Prozessor des ersten Rechnerplatzes eine neue Transaktion beginnt;
 15 Lesen der dem Aufzeichnungsschlüssel aus dem Speicher entsprechenden Zeitmarke dann, wenn der Prozessor des ersten Rechnerplatzes eine neue Transaktion beginnt;
 20 Vergleich einer dem Aufzeichnungsschlüssel entsprechenden bestehenden Zeitmarke mit der Zeitmarke, die der neuen Transaktion zugeordnet ist;
 25 Abbruch der neuen Transaktion dann, wenn die Zeitmarke, die der neuen Transaktion zugeordnet ist, kleiner als die dem Aufzeichnungsschlüssel entsprechende Zeitmarke ist; gekennzeichnet durch
 30 Erhöhung des Abbruchzählers für den ersten Rechnerplatz dann, wenn die neue Transaktion abgebrochen wird, um die Anzahl anzuzeigen, wie oft eine im ersten Rechnerplatz begonnene neue Transaktion abgebrochen wurde;
 35 Zugriff auf einen Abbruchzähler eines zweiten Rechnerplatzes;
 Vergleich des Abbruchzählers für den ersten Rechnerplatz, mit dem Abbruchzähler für den zweiten Rechnerplatz; und
 Änderung des gegenwärtigen Zeittakts für den ersten Rechnerplatz zur besseren Anpassung an den gegenwärtigen Zeittakt für den zweiten Rechnerplatz dann, wenn der Vergleich zwischen den Abbruchzählern für den ersten und zweiten Rechnerplatz zeigt, daß neue Transaktionen in einem der beiden Rechenplätze öfter abgebrochen wurden.

5

10

15

20

25

30

35

40

45

2. Verfahren nach Anspruch 1, bei dem:

der Schritt der Änderung des gegenwärtigen Zeittakts eine Änderung des gegenwärtigen Zeittakts um einen Betrag aufweist, der proportional zum Unterschied zwischen der Anzahl von Malen ist, in denen neue Transaktionen in den beiden Rechnerplätzen abgebrochen wurden.

50

3. Verfahren nach Anspruch 1, weiterhin umfassend:

Änderung des Abbruchzählers für den ersten Rechnerplatz zur besseren Anpassung an den Abbruchzähler für den zweiten Rechnerplatz dann, wenn der gegenwärtige Zeittakt für den ersten Rechnerplatz an den gegenwärtigen Zeittakt für den zweiten Rechnerplatz angepaßt wurde.

4. Verfahren nach Anspruch 1, bei dem

in dem Schritt der Änderung des gegenwärtigen Zeittakts nur dann eine Änderung des gegenwärtigen Zeittakts durchgeführt wird, wenn die Anzahl der im ersten Rechnerplatz abgebrochenen neuen Transaktionen die Anzahl der im zweiten Rechnerplatz abgebrochenen neuen Transaktionen übersteigt.

5. Verfahren nach Anspruch 1, bei dem

in dem Schritt der Änderung des gegenwärtigen Zeittakts nur dann eine Änderung des gegenwärtigen Zeittakts durchgeführt wird, wenn die Anzahl der im ersten Rechnerplatz abgebrochenen neuen Transaktionen die Anzahl der im zweiten Rechnerplatz abgebrochenen neuen Transaktionen um einen Schwellenwert übersteigt.

6. Verfahren nach Anspruch 1, weiterhin umfassend:

Eingabe eines neuen Lesezeittakts in der dem Aufzeichnungsschlüssel entsprechenden Zeitmarke, um anzuzeigen, wann auf den entsprechenden Dateneintrag durch den Prozessor zugegriffen wurde, wobei der neue Lesezeittakt auf den gegenwärtigen Zeittakt für den ersten Rechnerplatz bezogen ist.

7. Verfahren nach Anspruch 6, wobei jede Zeitmarke entsprechend dem Lesezeittakt eine Marke aufweist, die anzeigt, welcher Prozessor auf den entsprechenden Dateneintrag unter dem entsprechenden Aufzeichnungsschlüssel zugegriffen hat, und weiterhin umfassend:

Eingabe einer Marke in der dem Aufzeichnungsschlüssel entsprechenden Zeitmarke, wenn ein neuer Lesezeittakt in die Zeitmarke eingegeben wird, um anzuzeigen, ob auf den entsprechenden Dateneintrag durch den Prozessor in dem ersten Rechnerplatz zugegriffen wurde.

8. Verfahren nach Anspruch 1, bei dem der erste Rechnerplatz die Schritte des Zugriffs auf den Abbruchzähler für den zweiten Rechnerplatz, den Vergleich des Abbruchzählers für den ersten Rechnerplatz mit

dem Abbruchzähler für den zweiten Rechnerplatz und der Änderung des gegenwärtigen Zeittakts für den ersten Rechnerplatz zur besseren Anpassung an den gegenwärtigen Zeittakt für den zweiten Rechnerplatz wiederholt ausführt. 5

9. Verfahren nach Anspruch 1, bei dem der erste Rechnerplatz Eingabemittel (11) zur Eingabe einer Anforderung zur Taktanpassung durch den Bediener aufweist und bei dem der erste Rechnerplatz die Schritte des Zugriffs auf den Abbruchzählern für den zweiten Rechnerplatz, des Vergleichs des Abbruchzählers für den ersten Rechnerplatz mit dem Abbruchzähler für den zweiten Rechnerplatz und der Änderung des gegenwärtigen Zeittakts für den zweiten Rechnerplatz als Reaktion auf den Eingang einer Anforderung zur Taktanpassung durch den Bediener ausführt. 10

10. Verfahren, um die gegenwärtigen Zeittakte in einem verteilten Rechnersystem einzustellen, wobei das verteilte Rechnersystem eine Mehrzahl von Rechenplätzen (1, 2, 3) und jeder Rechnerplatz einen Prozessor (13), einen Speicher (12) mit einer Mehrzahl von Aufzeichnungsschlüsseln zur Speicherung einer zeitmarkentabelle, eine Datenbank (17), Erfolgsregistermittel zur Speicherung eines Erfolgszählers für den Rechnerplatz, Taktmittel (15) zur Erzeugung eines Taktsignals, das den Zeitablauf anzeigt, und Taktregistermittel (16) aufweist, die auf das Taktignal zur Speicherung eines gegenwärtigen Zeittakts für den Rechnerplatz reagieren, wobei die Datenbank jedes Rechnerplatzes eine Mehrzahl von Dateneinträgen aufweist, die der Mehrzahl von Aufzeichnungsschlüsseln entsprechen und auf die mittels eines Prozessors zugegriffen werden kann, wobei jeder Aufzeichnungsschlüssel einer Zeitmarke entspricht, die mindestens einen Zeittakt beinhaltet, der anzeigt, wann auf den entsprechenden Dateneintrag unter dem entsprechenden Aufzeichnungsschlüssel zuletzt zugegriffen wurde, wobei das verteilte Rechnersystem Netzwerkmittel (4) zur Verbindung der Mehrzahl von Rechenplätzen aufweist, um es einem Prozessor in einem Rechnerplatz zu ermöglichen, Dateneinträge zu senden und den Zugriff auf andere Rechnerplätze anzufordern, und wobei die folgenden Schritte durch einen Prozessor zumindestens in einem ersten der Rechnerplätze ausgeführt werden: 20

Zuteilung einer auf den gegenwärtigen Zeittakt des ersten Rechnerplatzes bezogenen Zeitmarke an eine neue Transaktion dann, wenn der Prozessor des ersten Rechnerplatzes eine neue Transaktion beginnt; 25

Lesen der dem Aufzeichnungsschlüssel aus dem Speicher entsprechenden Zeitmarke dann, wenn der Prozessor des ersten Rechner- 30

platzes eine neue Transaktion beginnt; Vergleich einer dem Aufzeichnungsschlüssel entsprechenden bestehenden Zeitmarke mit der Zeitmarke, die der neuen Transaktion zugeordnet ist; Abbruch der neuen Transaktion dann, wenn die Zeitmarke, die der neuen Transaktion zugeordnet ist, kleiner als die dem Aufzeichnungsschlüssel entsprechende Zeitmarke ist; gekennzeichnet durch Erhöhung des Erfolgszählers für den ersten Rechnerplatz dann, wenn die neue Transaktion nicht abgebrochen wurde, um die Anzahl anzuzeigen, wie oft eine im ersten Rechnerplatz begonnene neue Transaktion erfolgreich beendet wurde; Zugriff auf einen Erfolgszähler für einen zweiten Rechnerplatz; Vergleich des Erfolgszählers für den ersten Rechnerplatz mit dem Erfolgszähler für den zweiten Rechnerplatz; und Änderung des gegenwärtigen Zeittakts für den ersten Rechnerplatz zur besseren Anpassung an den gegenwärtigen Zeittakt für den zweiten Rechnerplatz dann, wenn der Vergleich zwischen den Erfolgszählern für den ersten und zweiten Rechnerplatz zeigt, daß neue Transaktionen in einem der beiden Rechenplätze öfter erfolgreich abgeschlossen wurden. 35

11. Verfahren nach Anspruch 10, bei dem der Schritt der Änderung des gegenwärtigen Zeittakts eine Änderung des gegenwärtigen Zeittakts um einen Betrag aufweist, der proportional zum Unterschied zwischen der Anzahl von Malen ist, in denen neue Transaktionen in den beiden Rechenplätzen erfolgreich beendet wurden. 40

12. Verfahren nach Anspruch 10, weiterhin umfassend: Änderung des Erfolgszählers für den ersten Rechnerplatz zur besseren Anpassung an den Erfolgszähler für den zweiten Rechnerplatz dann, wenn der gegenwärtige Zeittakt für den ersten Rechnerplatz an den gegenwärtigen Zeittakt für den zweiten Rechnerplatz angepaßt wurde. 45

13. Verfahren nach Anspruch 10, bei dem bei dem Schritt der Änderung des gegenwärtigen Zeittakts nur dann eine Änderung des gegenwärtigen Zeittakts durchgeführt wird, wenn die Anzahl der im ersten Rechnerplatz erfolgreich beendeten neuen Transaktionen die Anzahl der im zweiten Rechnerplatz erfolgreich 50

55

beendeten neuen Transaktionen übersteigt.	5	bediener ausführt.
14. Verfahren nach Anspruch 10, bei dem		Revendications
bei dem Schritt der Änderung des gegenwärtigen Zeittaks nur dann eine Änderung des gegenwärtigen Zeittaks durchgeführt wird, wenn die Anzahl der im ersten Rechnerplatz erfolgreich beendeten neuen Transaktionen die Anzahl der im zweiten Rechnerplatz erfolgreich beendeten neuen Transaktionen um einen Schwellenwert übersteigt.	10	1. Méthode d'ajustement des horocomptages actuels dans un système d'ordinateurs répartis, dans laquelle méthode le système d'ordinateurs répartis comporte une pluralité de sites d'ordinateur (1, 2, 3), chacun de ces sites d'ordinateur comprenant un processeur (13), une mémoire (12) ayant une pluralité de clés d'enregistrement pour stocker une table d'horodatage, une base de données (17), des moyens à registres de suspensions de transactions (14) pour stocker le comptage de suspension du site d'ordinateur, un moyen à horloge (15) pour générer un signal d'horloge indiquant l'écoulement du temps et un moyen à registre d'horloge (16) qui répond au signal d'horloge pour stocker l'horocomptage actuel du site d'ordinateur, dans laquelle méthode la base de données de chaque site d'ordinateur comporte une pluralité d'entrées de données (ou enregistrements) qui correspond à la pluralité de clés d'enregistrements et auxquelles les processeurs peuvent accéder, dans laquelle méthode chaque clé d'enregistrement correspond à un horodatage qui contient au moins un horocomptage qui indique l'instant auquel un processeur a accédé le plus récemment à l'entrée de données correspondante sous la clé d'enregistrement correspondante, dans laquelle méthode le système d'ordinateurs répartis comporte un moyen à réseau (4) pour interconnecter la pluralité de sites d'ordinateurs pour permettre au processeur d'un site d'ordinateur d'envoyer des entrées de données et demander le droit d'accéder à un autre site d'ordinateur, et dans laquelle méthode les phases suivantes sont exécutées par le processeur d'au moins un des sites d'ordinateur:
15. Verfahren nach Anspruch 10, weiterhin umfassend:	15	
Eingabe eines neuen Lesezeittaks in der dem Aufzeichnungsschlüssel entsprechenden Zeitmarke, um anzuzeigen, wann auf den entsprechenden Dateneintrag durch den Prozessor zugegriffen wurde, wobei der neue Lesezeitakt auf den gegenwärtigen Zeittakt für den ersten Rechnerplatz bezogen ist.	20	
16. Verfahren nach Anspruch 10, wobei jede Zeitmarke entsprechend dem Lesezeittakt eine Marke aufweist, die anzeigt, welcher Prozessor auf den entsprechenden Dateneintrag unter dem entsprechenden Aufzeichnungsschlüssel zugegriffen hat, und weiterhin umfassend:	25	
Eingabe einer Marke in die dem Aufzeichnungsschlüssel entsprechenden Zeitmarke, um anzuzeigen, daß auf den entsprechenden Dateneintrag durch den Prozessor in dem ersten Rechnerplatz zugegriffen wurde, wenn ein neuer Lesezeittakt in die Zeitmarke eingegeben wird.	30	
17. Verfahren nach Anspruch 10, bei dem der erste Rechnerplatz die Schritte des Zugriffs auf den Erfolgszähler für den zweiten Rechnerplatz, des Vergleichs des Erfolgszählers für den ersten Rechnerplatz mit dem Erfolgszähler für den zweiten Rechnerplatz und der Änderung des gegenwärtigen Zeittaks für den ersten Rechnerplatz zur besseren Anpassung an den gegenwärtigen Zeittakt für den zweiten Rechnerplatz wiederholt ausführt.	35	
18. Verfahren nach Anspruch 10, bei dem der erste Rechnerplatz Eingabemittel (11) zur Eingabe einer Anforderung zur Taktanpassung durch den Bediener aufweist und bei dem der erste Rechnerplatz die Schritte des Zugriffs auf den Erfolgszähler für den zweiten Rechnerplatz, des Vergleiches des Erfolgszählers für den ersten Rechnerplatz mit dem Erfolgszähler für den zweiten Rechnerplatz und der Änderung des gegenwärtigen Zeittaks für den zweiten Rechnerplatz als Reaktion auf den Eingang einer Anforderung zur Taktanpassung durch den	40	
	45	l'affectation d'un horodatage à une nouvelle transaction, cet horodatage étant fonction de l'horocomptage actuel du premier site d'ordinateur, lorsque le processeur du premier site d'ordinateur initie la nouvelle transaction;
	50	la lecture de l'horodatage de la clé d'enregistrement en mémoire, lorsque le processeur du premier site d'ordinateur initie la nouvelle transaction;
	55	la comparaison entre l'horodatage existant de la clé d'enregistrement et l'horodatage affecté à la nouvelle transaction;
		la suspension d'une nouvelle transaction lorsque l'horodatage affecté à cette nouvelle transaction est antérieur à l'horodatage existant de la clé d'enregistrement;
		caractérisée par une incrémentation du compteur de suspensions du premier site d'ordinateur lorsque la nouvelle transaction est suspen-

due, pour indiquer le nombre de fois que de nouvelles transactions initiées par le premier site d'ordinateurs ont été suspendues; l'accès au contenu du compteur de suspensions du second site d'ordinateur; la comparaison du contenu du compteur de suspension du premier site d'ordinateur au contenu du compteur de suspensions du second site d'ordinateur, et la modification de l'horocomptage actuel du premier site d'ordinateur pour l'aligner plus étroitement sur l'horocomptage actuel du second site d'ordinateur, lorsque la comparaison entre les contenus des compteurs de suspensions du premier et du second site d'ordinateur indique que les nouvelles transactions ont été suspendues plus fréquemment dans un des deux sites d'ordinateur. 5

2. Méthode selon la revendication 1, dans laquelle: 20

la phase de modification de l'horocomptage comprend la modification de l'horocomptage actuel d'une quantité de temps proportionnelle à la différence entre le nombre de fois que les nouvelles transactions ont été suspendues dans les deux sites d'ordinateur. 25

3. Méthode selon la revendication 1 et comprenant en outre: 30

la modification du contenu du compteur de suspensions du premier site d'ordinateur pour l'aligner plus étroitement sur le contenu du compteur de suspensions du second site d'ordinateur, lorsque l'horocomptage actuel du premier site d'ordinateur est modifié pour l'aligner plus étroitement sur l'horocomptage du second site d'ordinateur. 35

4. Méthode selon la revendication 1, dans laquelle: 40

la phase de modification de l'horocomptage actuel comprend la modification de l'horocomptage actuel que si le nombre de nouvelles transactions suspendues par le premier site d'ordinateur dépasse le nombre de nouvelles transactions suspendues par le second site d'ordinateur. 45

5. Méthode selon la revendication 1, dans laquelle la phase de modification de l'horocomptage comprend le changement de l'horocomptage uniquement lorsque le nombre de nouvelles transactions suspendues par le premier site d'ordinateur dépasse le nombre de nouvelles transactions suspendues par le second site d'ordinateur d'une quantité égale à une valeur-seuil. 50

6. Méthode selon la revendication 1 et comprenant en outre: 55

l'entrée d'un nouvel horocomptage lu dans la table des horodatages correspondant à la clé d'enregistrement pour indiquer lorsque le processeur a accédé aux données correspondantes, dans laquelle méthode le nouvel horodatage lu est fonction de l'horocomptage du premier site d'ordinateur.

7. Méthode selon la revendication 6, dans laquelle chaque horodatage contient un identifiant correspondant à chaque lecture d'horocomptage qui identifie le processeur qui a accédé à l'entrée de donnée (enregistrement) sous la clé d'enregistrement correspondante et qui comprend en outre:

l'entrée d'un identifiant dans la table d'horodatage correspondant à la clé d'enregistrement pour indiquer que le processeur du premier site d'ordinateur a accédé aux données correspondantes, lorsqu'une nouvelle lecture d'horocomptage est entrée dans la table d'horodatage.

8. Méthode selon la revendication 1, dans laquelle le premier site d'ordinateur exécute de manière récursive la phase d'accès au contenu du compteur des suspensions du second site d'ordinateur, la phase de comparaison du contenu du compteur des suspensions du premier site d'ordinateur au contenu du compteur des suspensions du second site d'ordinateur et la phase de modification de l'horocomptage du premier site d'ordinateur pour l'aligner plus étroitement à l'horocomptage actuel du second site d'ordinateur.

9. Méthode selon la revendication 1, dans laquelle le premier site d'ordinateur comporte un moyen d'entrée d'informations (11) pour entrer une demande d'alignement d'horloge par l'utilisateur, et dans laquelle le premier site d'ordinateur exécute la phase d'accès au contenu du compteur de suspensions du second site d'ordinateur et la phase de comparaison du contenu des compteurs de suspensions du premier site d'ordinateur et du second site d'ordinateur et la phase de modification de l'horocomptage actuel du premier site d'ordinateur pour l'aligner plus étroitement sur l'horocomptage du second site d'ordinateur en réponse à l'entrée d'une demande d'alignement d'horloge par l'utilisateur.

10. Méthode d'ajustement des horocomptages actuels dans un système d'ordinateurs répartis, dans laquelle méthode le système d'ordinateurs répartis comporte une pluralité de sites d'ordinateurs (1, 2, 3), chacun de ces sites d'ordinateur comprenant un

processeur (3), une mémoire (12) ayant une pluralité de clés d'enregistrement pour stocker une table d'horodatages, une base de données (17) un moyen à registre des transactions validées pour stocker le contenu du compteur de transactions validées du site d'ordinateur, un moyen à horloge (15) pour générer un signal d'horloge indiquant l'écoulement du temps, et un moyen à registre d'horloge (16) qui répond au signal d'horloge pour stocker l'horocomptage actuel du premier site, dans laquelle méthode la base de données de chaque site d'ordinateur comporte une pluralité d'entrées de données correspondant à la pluralité de clés d'enregistrement et auxquelles peuvent accéder les processeurs, dans laquelle méthode chaque clé d'enregistrement correspond à un horodatage qui contient au moins un horocomptage qui indique l'instant le plus récent où le processeur a accédé à l'entrée de données correspondante sous la clé d'enregistrement correspondante, dans laquelle méthode le système d'ordinateurs répartis comporte un moyen à réseau (14) pour interconnecter la pluralité de sites d'ordinateur pour permettre au processeur d'un site d'ordinateur d'envoyer des entrées de données et de demander des autorisations d'accès à un autre site d'ordinateur, et dans laquelle méthode les phases suivantes sont exécutées par le processeur d'au moins un des sites d'ordinateur:

l'affectation d'un horodatage à une nouvelle transaction, cet horodatage étant fonction de l'horocomptage du premier site d'ordinateur, lorsque le processeur du premier site d'ordinateur initie la nouvelle transaction;
 la lecture de l'horodatage correspondant à la clé d'enregistrement dans la mémoire, lorsque le processeur du premier site d'ordinateur initie la nouvelle transaction;
 la comparaison entre l'horodatage existant de la clé d'enregistrement et l'horodatage affecté à la nouvelle transaction;
 la suspension de la nouvelle transaction, lorsque l'horodatage affecté à la nouvelle transaction est antérieur à l'horodatage existant de la clé d'enregistrement, 45 caractérisée par l'incrémentation du compteur de transactions validées du premier site d'ordinateur, lorsque la nouvelle transaction est validée, pour indiquer le nombre de fois que les nouvelles transactions initiées par le premier site d'ordinateur ont été accomplies avec succès;
 l'accès au contenu du compteur de transactions validées du second site d'ordinateur;
 la comparaison entre le contenu du compteur de transactions validées du premier site d'ordinateur et le contenu du compteur des suspensions du second site d'ordinateur, et

5 la modification de l'horocomptage actuel du premier site d'ordinateur pour qu'il s'aligne plus étroitement sur l'horocomptage du second site d'ordinateur lorsque la comparaison entre les contenus des compteurs de transactions validées du premier et du second site d'ordinateur indique que des nouvelles transactions ont été effectuées avec succès plus fréquemment dans un des deux sites d'ordinateurs.

10 11. Méthode selon la revendication 10, dans laquelle:

15 la phase de modification de l'horocomptage actuel comprend la modification de l'horocomptage actuel d'une quantité de temps qui est proportionnelle à la différence entre le nombre de fois que des transactions ont été validées dans un des deux sites d'ordinateurs.

20 12. Méthode selon la revendication 10 et comprenant en outre:

25 la modification du contenu du compteur de transactions validées du premier site d'ordinateur pour l'aligner plus étroitement sur le contenu du compteur de transactions validées du second site d'ordinateur, lorsque l'horocomptage actuel du premier site d'ordinateur est modifié pour l'aligner plus étroitement sur l'horocomptage actuel du second site d'ordinateur.

30 13. Méthode selon la revendication 10, dans laquelle:

35 la phase de modification de l'horocomptage actuel comprend le changement de l'horocomptage actuel uniquement si le nombre de nouvelles transactions validées dans le premier site d'ordinateur dépasse le nombre de nouvelles transactions validées dans le second site d'ordinateur.

40 14. . Méthode selon la revendication 10, dans laquelle:

45 la phase de modification de l'horocomptage actuel comprend la modification de l'horocomptage uniquement lorsque le nombre de nouvelles transactions validées dans le premier site d'ordinateur dépasse le nombre de nouvelles transactions validées dans le second site d'ordinateur d'une valeur-seuil déterminée.

50 15. Méthode selon la revendication 10, qui comprend en outre:

55 l'entrée d'une nouvelle lecture d'horocomptage de la clé d'enregistrement correspondant à la clé d'enregistrement pour indiquer l'accès le plus récent de l'entrée de données par le pro-

cesseur, dans laquelle méthode la nouvelle lecture d'horocomptage est fonction de l'horocomptage actuel du premier site d'ordinateur.

16. Méthode selon la revendication 10, dans laquelle chaque horodatage contient un identifiant correspondant à chaque lecture d'horocomptage qui identifie le processeur qui a accédé à l'entrée de données (enregistrement) sous la clé d'enregistrement correspondante, et qui comprend en outre: 5

l'entrée d'un identifiant dans l'horodatage correspondant à la clé d'enregistrement pour indiquer que le processeur du premier site d'ordinateur a accédé à l'entrée de données correspondante, lorsqu'une nouvelle lecture d'horocomptage a été mise dans l'horodatage. 10 15

17. Méthode selon la revendication 10, dans laquelle le premier site d'ordinateur exécute de manière récursive les phases d'accès au contenu du compteur de transactions validées du second site d'ordinateur, de comparaison du contenu du compteur de transactions validées pour le premier site d'ordinateur et du contenu du compteur de transactions validées du second site d'ordinateur et de changement de l'horocomptage actuel du premier site d'ordinateur pour l'aligner plus étroitement sur le second site d'ordinateur. 20 25

18. Méthode selon la revendication 10, dans laquelle le premier site d'ordinateur comporte un moyen d'entrée d'informations (11) pour entrer une demande d'alignement d'horloge par l'utilisateur, et dans laquelle méthode le premier site d'ordinateur exécute la phase d'accès au contenu du compteur de transactions validées du premier site d'ordinateur avec le contenu du compteur de transactions validées du second site d'ordinateur, et la phase de modification de l'horocomptage actuel du premier site d'ordinateur pour l'aligner plus étroitement sur l'horocomptage actuel du second site d'ordinateur, en réponse à une entrée d'une demande d'alignement d'horloge par l'utilisateur. 30 35 40 45

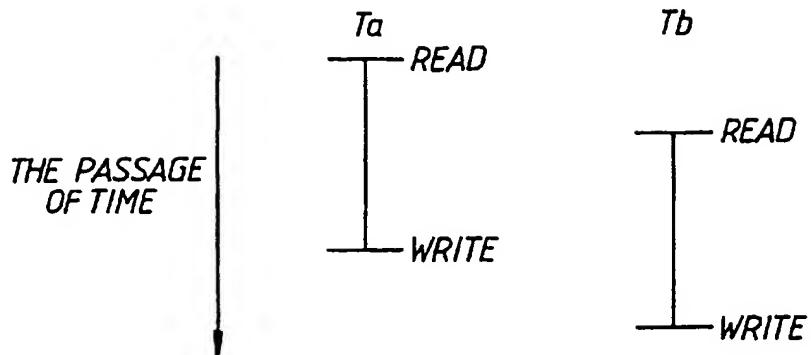


Fig. 1. (PRIOR ART)

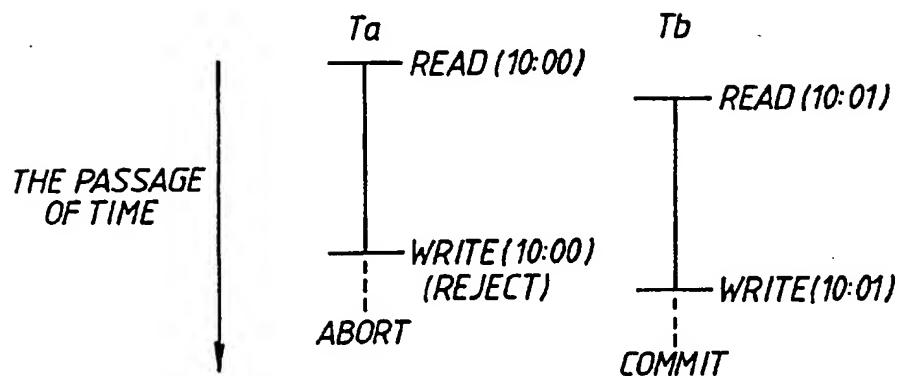


Fig. 3A (PRIOR ART)

RECORD KEY	READ TIMESTAMP	WRITE TIMESTAMP
100	0:00	0:00
		↓
100	10:00(A)	0:00
		↓
100	10:01(B)	0:00

Fig. 3B (PRIOR ART)

OPERATION	READ
RECORD KEY	100
TABLE NAME	ACCOUNT

Fig. 2.

TIMESTAMP TABLE				DATABASE (TABLE ACCOUNT)			
RECORD KEY	READ TIMESTAMP	WRITE TIMESTAMP		KEY	NAME	BALANCE	DEPOSIT
100	10:00 (A)	10:00 (A)		100	MIKE	\$ 1,000	
101	11:00 (B)	10:00 (C)		101	JANE	\$ 2,000	
102	9:00 (A)	7:00 (B)		102	BOB	\$ 3,000	
103	9:00 (B)	11:00 (C)		*	*	*	*
*	*	*		*	*	*	*
*	*	*		*	*	*	*

Fig. 4.

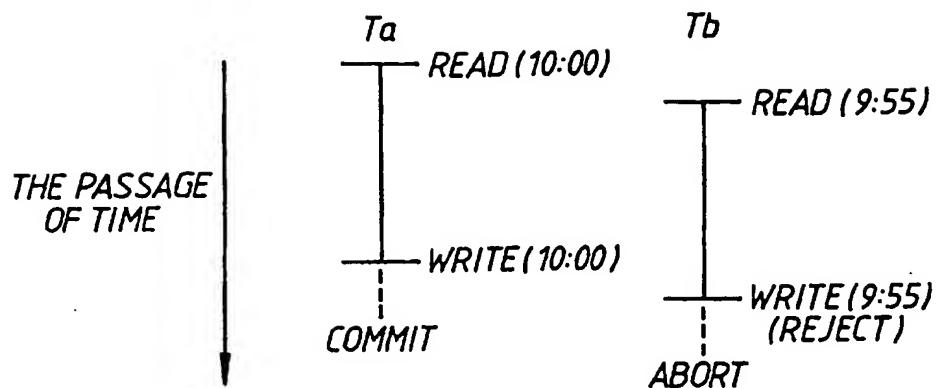


Fig. 5A (PRIOR ART)

↓

RECORD KEY	READ TIMESTAMP	WRITE TIMESTAMP
100	10:00(A)	0:00

↓

RECORD KEY	READ TIMESTAMP	WRITE TIMESTAMP
100	9:55(B)	0:00

↓

RECORD KEY	READ TIMESTAMP	WRITE TIMESTAMP
100	9:55(B)	10:00(A)

Fig. 5B (PRIOR ART)

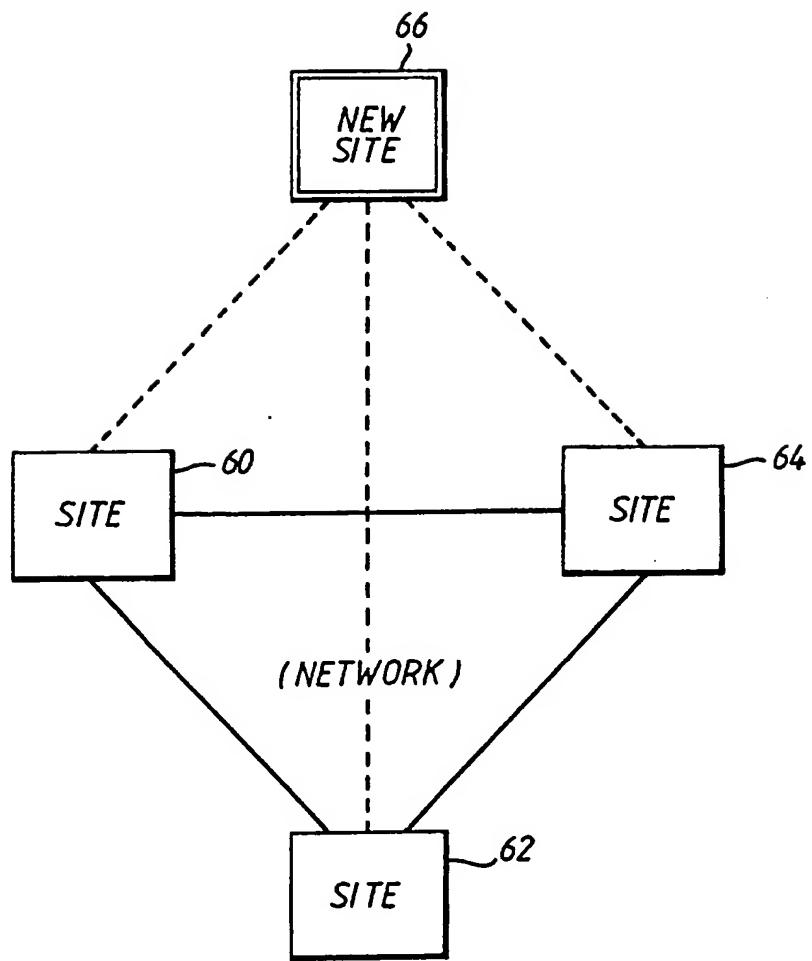


Fig. 6. (PRIOR ART)

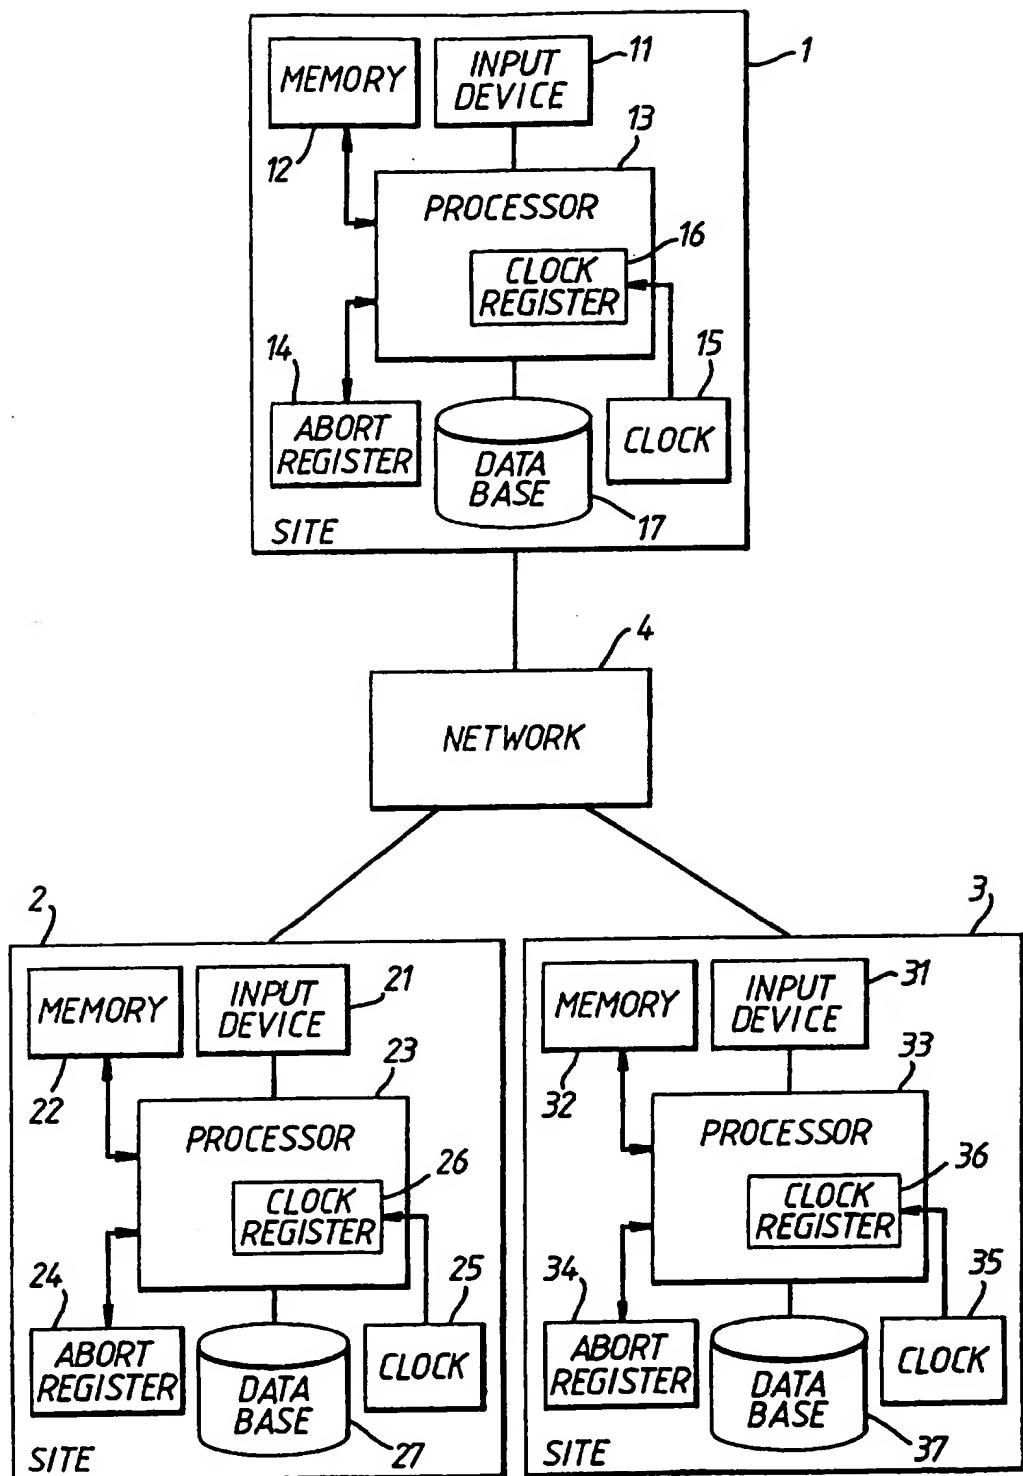


Fig. 7.

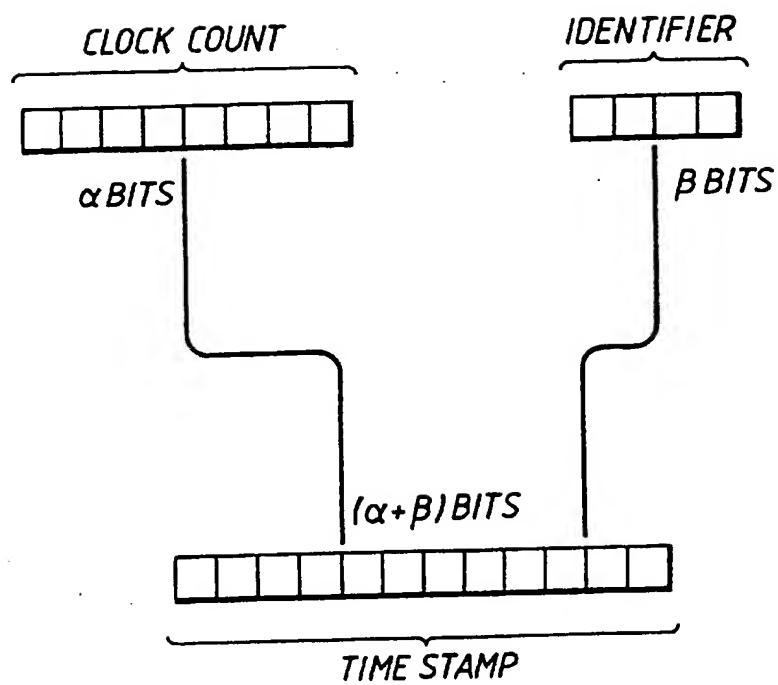


Fig.8.

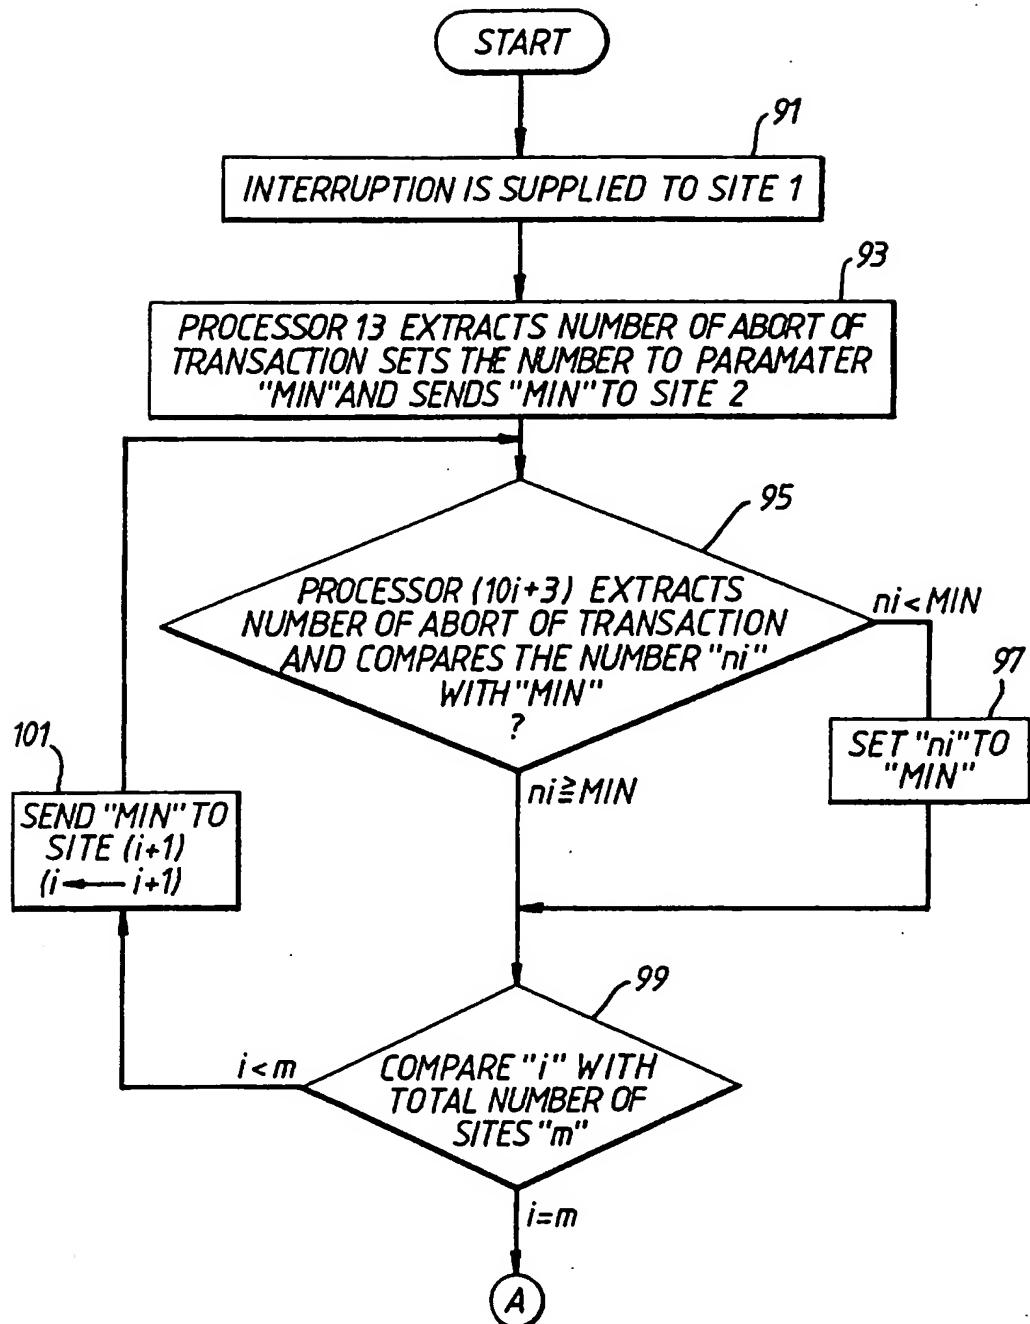


Fig. 9A

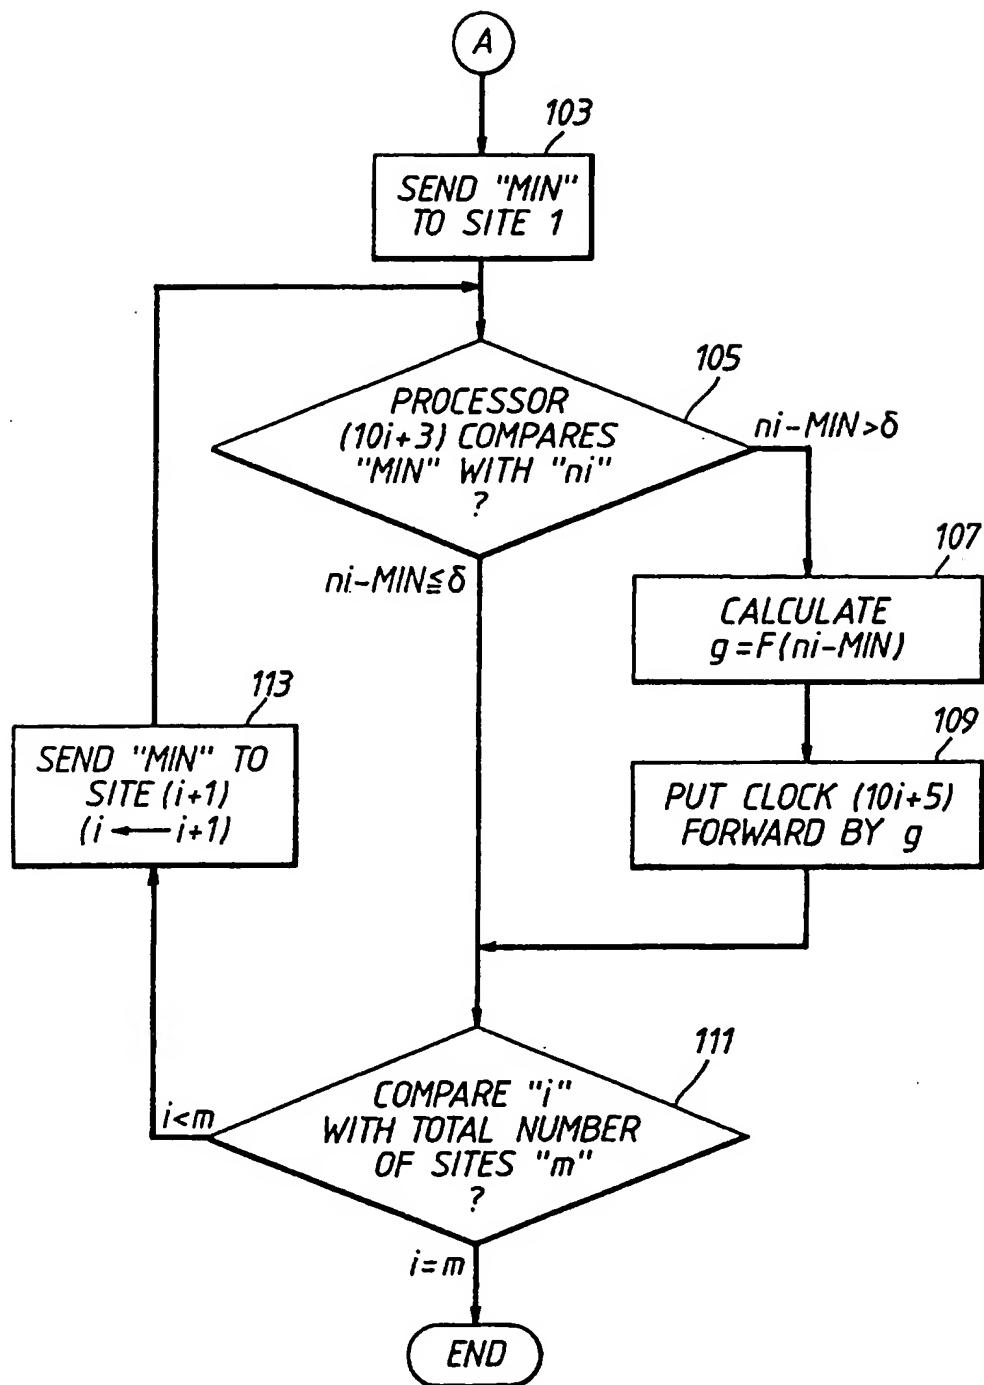


Fig.9B